



Microsoft

70-475 Exam

Microsoft Designing and Implementing Big Data Analytics Solutions Exam

Thank you for Downloading 70-475 exam PDF Demo

You can Buy Latest 70-475 Full Version Download

<https://www.certkillers.net/Exam/70-475>

<https://www.certkillers.net>

Version: 13.0

Case Study: 1

Relecloud General Overview

Relecloud is a social media company that processes hundreds of millions of social media posts per day and sells advertisements to several hundred companies.

Relecloud has a Microsoft SQL Server database named DB1 that stores information about the advertisers. DB1 is hosted on a Microsoft Azure virtual machine.

Physical locations

Relecloud has two main offices. The offices are located in San Francisco and New York City.

The offices are connected to each other by using a site-to-site VPN. Each office connects directly to the Internet.

Business model

Relecloud modifies the pricing of its advertisements based on trending topics. Topics are considered to be trending if they generate many mentions in a specific country during a 15-minute time frame. The highest trending topics generate the highest advertising revenue.

CTO statement

Relecloud wants to deliver reports to the advertisers by using Microsoft Power BI. The reports will provide real-time data on trending topics, current advertising rates, and advertising costs for a given month.

Relecloud will analyze the trending topics data, and then store the data in a new data warehouse for ad-hoc analysis. The data warehouse is expected to grow at a rate of 1 GB per hour or 8.7 terabytes (TB) per year. The data will be retained for five years for the purpose of long term trending.

Requirements

Business goals

Management at Relecloud must be able to view which topics are trending to adjust advertising rates in near real-time.

Planned changes

Relecloud plans to implement a new streaming analytics platform that will report on trending topics. Relecloud plans to implement a data warehouse named DB2.

General technical requirements

Relecloud identifies the following technical requirements:

- Social media data must be analyzed to identify trending topics in real time.
- The use of Infrastructure as a Service (IaaS) platforms must be minimized, whenever possible.
- The real-time solution used to analyze the social media data must support scaling up and down without service interruption.

Technical requirements for advertisers

Relecloud identifies the following technical requirements for the advertisers

- The advertisers must be able to see only their own data in the Power BI reports.
- The advertisers must authenticate to Power BI by using Azure Active Directory (Azure AD) credentials.
- The advertisers must be able to leverage existing Transact-SQL language knowledge when developing the real-time streaming solution.
- Members of the internal advertising sales team at Relecloud must be able to see only the sales data of the advertisers to which they are assigned.
- The Internal Relecloud advertising sales team must be prevented from inserting, updating, and deleting rows for the advertisers to which they are not assigned.
- The internal Relecloud advertising sales team must be able to use a text file to update the list of advertisers, and then to upload the file to Azure Blob storage.

DB1 requirements

Relecloud identifies the following requirements for DB1:

- Data generated by the streaming analytics platform must be stored in DB1.
- The user names of the advertisers must be mapped to CustomerID in a table named Table2.
- The advertisers in DB1 must be stored in a table named Table1 and

must be refreshed nightly.

- The user names of the employees at Relecloud must be mapped to EmployeeID in a table named Table3.

DB2 requirements

Relecloud identifies the following requirements for DB2:

- DB2 must have minimal storage costs.
- DB2 must run load processes in parallel.
- DB2 must support massive parallel processing.
- DB2 must be able to store more than 40 TB of data.
- DB2 must support scaling up and down, as required.
- Data from DB1 must be archived in DB2 for long-term storage.
- All of the reports that are executed from DB2 must use aggregation.
- Users must be able to pause DB2 when the data warehouse is not in use.
- Users must be able to view previous versions of the data in DB2 by using aggregates.

ETL requirements

Relecloud identifies the following requirements for extract, transformation, and load (ETL):

- Data movement between DB1 and DB2 must occur each hour.
- An email alert must be generated when a failure of any type occurs during ETL processing.

rls_table1

You execute the following code for a table named rls_table1.

```

create function rls_table1 (@CustomerId int, @SalesPersonId int)
    returns table
    with schemabinding
as
return
select 1 as result
from dbo.table1
join dbo.table2 on table1.customerid = Table2.CustomerId
where table2.UserName = suser_sname()
    and table1.customerid = @CustomerId
union all
select 1 as result
from dbo.table1
join dbo.table3 on table1.salespersonid = table3.EmployeeId
where table3.UserName = suser_sname()
    and table1.salespersonid = @SalesPersonId
go

```

dbo.table1

You use the following code to create Table1.

```

create table table1
(customerid int,
salespersonid int
...
)
Go

```

Streaming data

The following is a sample of the Streaming data.

User	Country	Topic	Time
user1	USA	Topic1	2017-01-01T00:00:01.0000000Z
user1	USA	Topic3	2017-01-01T00:02:01.0000000Z
user2	Canada	Topic2	2017-01-01T00:01:11.0000000Z
user3	India	Topic1	2017-01-01T00:03:14.0000000Z

Question: 1

Which technology should you recommend to meet the technical requirement for analyzing the social media data?

- A. Azure Stream Analytics
- B. Azure Data Lake Analytics
- C. Azure Machine Learning
- D. Azure HDInsight Storm clusters

Answer: A

Azure Stream Analytics is a fully managed event-processing engine that lets you set up real-time analytic computations on streaming data.

Scalability

Stream Analytics can handle up to 1 GB of incoming data per second. Integration with Azure Event Hubs and Azure IoT Hub allows jobs to ingest millions of events per second coming from connected devices, clickstreams, and log files, to name a few. Using the partition feature of event hubs, you can partition computations into logical steps, each with the ability to be further partitioned to increase scalability.

Question: 2

DRAGDROP

You need to implement a solution that meets the data refresh requirement for DB1. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions**Answer Area**

In DB1, create external objects.

From the Azure portal, export the storage account key.

In DB1, create a stored procedure that imports data from an external table to Table1.

From the Azure portal, create and schedule an Azure Automation job that executes the stored procedure.

In DB1, create a staging table.



Answer:

Answer Area

In DB1, create a stored procedure that imports data from an external table to Table1.

From the Azure portal, export the storage account key.

From the Azure portal, create and schedule an Azure Automation job that executes the stored procedure.

Azure Data Factory can be used to orchestrate the execution of stored procedures. This allows more complex pipelines to be created and extends Azure Data Factory's ability to leverage the computational power of SQL Data Warehouse.

From scenario:

Relecloud has a Microsoft SQL Server database named DB1 that stores information about the advertisers. DB1 is hosted on a Microsoft Azure virtual machine.

Relecloud identifies the following requirements for DB1:

- Data generated by the streaming analytics platform must be stored in DB1.
- The advertisers in DB1 must be stored in a table named Table1 and must be refreshed nightly.

Question: 3

DRAG DROP

You need to create a query that identifies the trending topics.

How should you complete the query? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

- DATETIME
- GROUP BY
- HoppingWindow
- ORDER BY
- SlidingWindow
- TIMESTAMP
- VIEW BY

Answer Area

SELECT Country, Topic, count(*)

FROM Input **Value** BY Time

Value Country, Topic, **Value** (minute, 15)

Answer:

Answer Area

SELECT Country, Topic, count(*)

FROM Input **TIMESTAMP** BY Time

GROUP BY Country, Topic, **SlidingWindow** (minute, 15)

From scenario: Topics are considered to be trending if they generate many mentions in a specific country during a 15-minute time frame.

Box 1: TimeStamp

Azure Stream Analytics (ASA) is a cloud service that enables real-time processing over streams of data flowing in from devices, sensors, websites and other live systems. The stream-processing logic in ASA is expressed in a SQL-like query language with some added extensions such as windowing for performing temporal calculations.

ASA is a temporal system, so every event that flows through it has a timestamp. A timestamp is assigned automatically based on the event's arrival time to the input source but you can also access a timestamp in your event payload explicitly using **TIMESTAMP BY**:

```
SELECT * FROM SensorReadings TIMESTAMP BY time
```

Box 2: GROUP BY

Example: Generate an output event if the temperature is above 75 for a total of 5 seconds

```
SELECT sensorId, MIN(temp) as temp  
FROM SensorReadings  
TIMESTAMP BY time  
GROUP BY sensorId, SlidingWindow(second, 5)  
HAVING MIN(temp) > 75
```

Box 3: SlidingWindow

Windowing is a core requirement for stream processing applications to perform set-based operations like counts or aggregations over events that arrive within a specified period of time. ASA supports three types of windows: Tumbling, Hopping, and Sliding.

With a Sliding Window, the system is asked to logically consider all possible windows of a given length and output events for cases when the content of the window actually changes – that is, when an event entered or existed the window.

Question: 4

HOTSPOT

You implement DB2.

You need to configure the tables in DB2 to host the data from DB1. The solution must meet the requirements for DB2.

Which type of table and history table storage should you use for the tables? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer area

Table:
Change Data Capture
Change tracking
Temporal table

History table storage:
Clustered columnstore
In-Memory OLTP
Row store

Answer:

Table:
Change Data Capture
Change tracking
Temporal table

History table storage:
Clustered columnstore
In-Memory OLTP
Row store

From Scenario: Relecloud plans to implement a data warehouse named DB2.

Box 1: Temporal table

From Scenario:

Relecloud identifies the following requirements for DB2:

Users must be able to view previous versions of the data in DB2 by using aggregates.

DB2 must be able to store more than 40 TB of data.

A system-versioned temporal table is a new type of user table in SQL Server 2017, designed to keep a full history of data changes and allow easy point in time analysis.

A temporal table also contains a reference to another table with a mirrored schema

a. The system uses this table to automatically store the previous version of the row each time a row in the temporal table gets updated or deleted. This additional table is referred to as the history table, while the main table that stores current (actual) row versions is referred to as the current table or simply as the temporal table.

Question: 5

DRAG DROP

You need to implement rls_table1.

Which code should you execute? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

● ● ● ● ●

Values	Answer Area
Block	CREATE [Value] POLICY dbo.rls_table1_policy
Filter	ADD [Value] PREDICATE dbo.rls_table1(CustomerId, salespersonid) on dbo.table1,
Grant	ADD [Value] PREDICATE dbo.rls_table1(CustomerId, salespersonid) on dbo.table1 BEFORE UPDATE,
Security	ADD [Value] PREDICATE dbo.rls_table1(CustomerId, salespersonid) on dbo.table1 BEFORE DELETE,
Server	ADD [Value] PREDICATE dbo.rls_table1(CustomerId, salespersonid) on dbo.table1 AFTER INSERT
	with (state = on)

Answer:

Answer Area

```

CREATE [ Security ] POLICY dbo.rls_table1_policy
ADD [ Filter ] PREDICATE dbo.rls_table1(CustomerId, salespersonid) on dbo.table1,
ADD [ Block ] PREDICATE dbo.rls_table1(CustomerId, salespersonid) on dbo.table1 BEFORE UPDATE,
ADD [ Block ] PREDICATE dbo.rls_table1(CustomerId, salespersonid) on dbo.table1 BEFORE DELETE,
ADD [ Filter ] PREDICATE dbo.rls_table1(CustomerId, salespersonid) on dbo.table1 AFTER INSERT
with ( state = on )
  
```

Box 1: Security

Security Policy

Example: After we have created Predicate function, we have to bind it to the table, using Security Policy. We will be using CREATE SECURITY POLICY command to set the security policy in place.

```
CREATE SECURITY POLICY DepartmentSecurityPolicy
ADD FILTER PREDICATE dbo.DepartmentPredicateFunction(UserDepartment) ON
dbo.Department
WITH(STATE = ON)
```

Box 2: Filter

[FILTER | BLOCK]

The type of security predicate for the function being bound to the target table. FILTER predicates silently filter the rows that are available to read operations. BLOCK predicates explicitly block write operations that violate the predicate function.

Box 3: Block

Box 4: Block

Box 5: Filter

Question: 6

HOTSPOT

Which service solution and which table storage solution should you recommend for DB2? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Service:

An Azure virtual machine that has SQL Server installed	▼
Azure SQL Data Warehouse	
Azure SQL Database	

Table storage:

Clustered columnstore index	▼
Clustered index	
In-Memory OLTP	

Answer:

Service:

An Azure virtual machine that has SQL Server installed	▼
Azure SQL Data Warehouse	
Azure SQL Database	

Table storage:

Clustered columnstore index	▼
Clustered index	
In-Memory OLTP	

Explanation:

Box 1: Azure SQL Data Warehouse

Scenario: Relecloud plans to implement a data warehouse named DB2.

Box 2: Clustered Columnstore index

Columnstore index is a new type of index introduced in SQL Server 2012. It is a column-based non-clustered index geared toward increasing query performance for workloads that involve large amounts of data, typically found in data warehouse fact tables.

A clustered columnstore index is the physical storage for the entire table.

Scenario:

Relecloud identifies the following requirements for DB2:

DB2 must be able to store more than 40 TB of data.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview>

Question: 7

HOTSPOT

You need to configure the alert to meet the requirements for ETL.

Which settings should you use for the alert? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Event:

	▼
Activity Run Finished	
Activity Run Started	
On-Demand HDI Cluster Create Start	
On-Demand HDI Cluster Created Successfully	
On-Demand HDI Cluster Deleted	

Status:

	▼
Failed	
Succeeded	

Substatus:

	▼
--	
Abandoned	
Failed Execution	
Failed Resource Allocation	
Failed Validation	
Timed Out	

Answer:

Event:

Activity Run Finished	▼
Activity Run Started	
On-Demand HDI Cluster Create Start	
On-Demand HDI Cluster Created Successfully	
On-Demand HDI Cluster Deleted	

Status:

Failed	▼
Succeeded	

Substatus:

--	▼
Abandoned	
Failed Execution	
Failed Resource Allocation	
Failed Validation	
Timed Out	

Explanation:

Scenario: Relecloud identifies the following requirements for extract, transformation, and load (ETL):
An email alert must be generated when a failure of any type occurs during ETL processing.

Case Study: 3

Litware, Inc

Overview

General Overview

Litware, Inc. is a company that manufactures personal devices to track physical activity and other health-related data.

Litware has a health tracking application that sends health-related data from a user's personal device to Microsoft Azure.

Physical Locations

Litware has three development and commercial offices. The offices are located in the United States, Luxembourg, and India.

Litware products are sold worldwide. Litware has commercial representatives in more than 80 countries.

Existing Environment

Environment

In addition to using desktop computers in all of the offices. Litware recently started using Microsoft Azure resources and services for both development and operations. Litware has an Azure Machine Learning Solution.

Litware Health Tracking Application

Litware recently extended its platform to provide third-party companies with the ability to upload data from devices to Azure. The data can be aggregated across multiple devices to provide users with a comprehensive view of their global health activity.

While the upload from each device is small, potentially more than 100 million devices will upload data daily by using an Azure event hub.

Each health activity has a small amount of data, such as activity type, start date/time, and end date/time. Each activity is limited to a total of 3 KB and includes a customer Identification key.

In addition to the Litware health tracking application, the users' activities can be reported to Azure by using an open API.

Machine Learning Experiments

The developers at Litware perform Machine Learning experiments to recommend an appropriate health activity based on the past three activities of a user.

The Litware developers train a model to recommend the best activity for a user based on the hour of the day.

Requirements**Planned Changes**

Litware plans to extend the existing dashboard features so that health activities can be compared between the users based on age, gender, and geographic region.

Business Goals

Minimize the costs associated with transferring data from the event hub to Azure Storage.

Technical Requirements

Litware identifies the following technical requirements:

Data from the devices must be stored from three years in a format that enables the fast processing of data fields and Filtering.

The third-party companies must be able to use the Litware Machine learning models to generate recommendations to their users by using a third-party application.

Any changes to the health tracking application must ensure that the Litware developers can run the experiments without interrupting or degrading the performance of the production environment.

Privacy Requirements

Activity tracking data must be available to all of the Litware developers for experimentation. The developers must be prevented from accessing the private information of the users.

Other Technical Requirements

When the Litware health tracking application asks users how they feel, their responses must be reported to Azure.

Question: 8

You need to recommend a data handling solution to support the planned changes to the dashboard. The solution must meet the privacy requirements.

What is the best recommendation to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

- A. anonymization
- B. encryption
- C. obfuscation
- D. compression

Answer: C

Question: 9

You need to recommend a data transfer solution to support the business goals. What should you recommend?

- A. Configure the health tracking application to cache data locally for 24 hours.
- B. Configure the health tracking application to Aggregate activities in blocks of 128 KB.
- C. Configure the health tracking application to cache data locally for 12 hours.
- D. Configure the health tracking application to aggregate activities in blocks of 64 KB.

Answer: D

Question: 10

You need to recommend a permanent Azure Storage solution for the activity data. The solution must meet the technical requirements.

What is the best recommendation to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Azure SQL Database
- B. Azure Queue storage
- C. Azure Blob storage
- D. Azure Event Hubs

Answer: A

Question: 11

Users report that when they access data that is more than one year old from a dashboard, the response time is slow.

You need to resolve the issue that causes the slow response when visualizing older data. What should you do?

- A. Process the event hub data first, and then process the older data on demand.
- B. Process the older data on demand first, and then process the event hub data.
- C. Aggregate the older data by time, and then save the aggregated data to reference data streams.
- D. Store all of the data from the event hub in a single partition.

Answer: C

Question: 12

You extend the dashboard of the health tracking application to summarize fields across several users.

You need to recommend a file format for the activity data in Azure that meets the technical requirements.

What is the best recommendation to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

- A. ORC
- B. TSV
- C. CSV
- D. JSON
- E. XML

Answer: E

Thank You for trying 70-475 PDF Demo

To Buy Latest 70-475 Full Version Download visit link below

<https://www.certkillers.net/Exam/70-475>

Start Your 70-475 Preparation

[Limited Time Offer] Use Coupon “CKNET” for Further discount on your purchase. Test your 70-475 preparation with actual exam questions.

<https://www.certkillers.net>